SPANNING TREE ALTERNATE ROUTING BRIDGE PROTOCOL

BACKGROUND OF THE INVENTION

FIELD OF THE INVENTION

The present invention relates generally to employing a bridge protocol for interconnecting two or more local area networks (LANs). More particularly the invention relates to apparatus and method, which is backward compatible with existing 802.1D Spanning Tree Bridge Protocol, for improving routing capability of spanning tree forwarding without a significant increase in complexity by providing a shorter alternate forwarding path if possible while using a path on the spanning tree by default.

BACKGROUND INFORMATION

A *Local Area Network (LAN)* is used to connect end stations together within close distance in order to provide high-bandwidth communications. A single LAN has a limited number of end stations, a limited size, and a limited amount of offered load. In this respect, LANs cannot grow beyond a certain limit. LANs may be interconnected via internetworking devices such as bridges and routers. These devices have different advantages and disadvantages depending on the internetworking environment. In the early days of internetworking, bridges were popular because they were much cheaper and faster than routers. In addition, bridges were used to support heterogeneous network layer protocols. The primitive computing technology of those days favored off-loading of work to larger servers using protocols that were optimized for LANs.

IEEE 802 Standards Committee has specified two bridge protocols. IEEE 802.1 group has issued the IEEE 802.1D Spanning Tree Bridge Protocol and IEEE 802.5 group has issued the Source Routing Bridge Protocol. Among these two schemes, IEEE 802.1D offers a better solution and has been studied more intensively. This approach is transparent to end stations and requires no modifications to the MAC layer of end stations. All the routing related operations are done in the bridges. Today, the IEEE 802.1D Spanning Tree Bridge Protocol is widely used for interconnecting the family of IEEE 802 standard LANs. For example, the Data-Over-Cable Service Interface Specifications (DOCSIS) describes the use of the IEEE 802.1D Spanning Tree Bridge Protocol to interconnect Cable Modem Termination Systems (CMTSs) over a switched or bridged headend network. According to DOCSIS, data forwarding through the CMTS may be transparent bridging, or network layer forwarding, but data forwarding through the Cable Modem (CM) is link-layer transparent bridging. The IEEE 802.1D standard is optional for CMs intended for residential use, but CMs intended for commercial use and bridging CMTSs must support the IEEE 802.1D standard.

A bridge has several *ports* connecting to different LANs. A frame sent from one LAN to the other will typically go through one or more ports and bridges. As bridges are capable of filtering frames, they are useful for dealing with unnecessary broadcast traffic. Such a broadcast containment capability renders bridging a simple solution to implementing a virtual LAN. This bridged LAN environment should be

transparent and looks like a single LAN to end users. The basic function of bridges is to forward MAC (Medium Access Control) frames from one LAN to another, therein providing an extension to the LAN without requiring any modification to the communications software in the end stations attached to the LANs. Bridges do not modify the content or format of the MAC frames they receive. The operation of bridges should not misorder or duplicate frames. Upper-layer protocol transparency is a primary advantage of bridging since bridges can rapidly forward traffic representing any network-layer protocol without having to examine upper-layer information.

The landscape for internetworking has evolved considerably with advances in high-speed layer 3 routing and layer 2 switching technologies. Functionalities at the two layers are increasingly similar. While routers are generally more intelligent than bridges in terms of their dynamic routing capability, they are also more complicated and costly to implement. Bridges have been designed to span a range of routing capabilities from dynamic source routing to static spanning tree forwarding, thereby allowing a trade-off between routing performance and protocol complexity. Although routers are becoming cheaper and faster than they used to be, they remain more complicated than bridges to operate because intermediate hops must still rise above layer 2. In spite of the common wisdom that IP has won the network layer, there are still going to be non-IP layer 3 protocols in the foreseeable future. On the other hand, while bridges are evolving to accommodate more and more layer 3 functionality, they will always support multiple layer 3 protocols.

An IP (Internet Protocol) address encodes both a network and a host on that network. Since it does not specify an individual machine, but a connection to a network, the IP address of a host must change whenever it moves from one network to another. On the other hand, an IEEE 802 MAC address identifies a physical interface from a station to a LAN, and is always applicable no matter where the station is plugged into a network. Such portability of end station addresses is important particularly for mobility and the benefit of plug-and-play. Although new features, are emerging to minimize the need to configure and reconfigure IP addresses, these features can increase the cost and processing overhead of the system. DHCP (Dynamic Host Configuration Protocol), for example, provides a widely deployed framework for host registration and configuration. DHCP, however, was designed only for fixed hosts on physically secure LANs. DHCP is being extended to allow dynamic reconfiguration of a single host triggered by the DHCP server (e.g. a new IP address). Depending on the bandwidth of the network between server and client, the delay in the reconfiguration process can grow exponentially as failed retransmissions trigger exponential backoff.

In the IEEE 802.1D standard, a shortest path spanning tree with respect to a predetermined bridge, known as a *root bridge*, is used to interconnect LANs to form an extended LAN. A frame sent from one LAN to another could follow a longer path on the spanning tree than necessary when there exists an alternative shorter path connecting them. Note that *non-tree links*, which are links that have not been selected by the 802.1D spanning tree algorithm, are not used to share the load of the traffic. The load around the root bridge may be heavy, and throughput is severely limited.

The IEEE 802.1D specification defines a protocol architecture for MAC bridges and recommends formats for a globally administered set of MAC station addresses across multiple LANs.

FIG. 1 shows a bridge protocol architecture for a connection of two LANs via local 10 or remote 12, 14 bridging. Referring to the OSI (open systems interconnect) reference model, a bridge encompasses the first two layers, namely the Physical Layer (layer 1) and the Data Link Layer (layer 2). There are two sublayers in layer 2: Medium Access Control (MAC) sublayer and Logical Link Control (LLC) sublayer. Bridges operate relay functions on the MAC sublayer and interface with the LLC sublayer above through LLC service access points. By using bridges, a growing LAN can be partitioned into self-contained units for administrative or maintenance reasons, as well as to improve performance via load balancing and fault isolation. Bridges are typically used to interconnect LANs of the same type, such as the family of IEEE 802 LANs. Translation among different link-layer protocols is needed, however, when the interconnected LANs are not homogeneous (e.g., IEEE 802.3 and IEEE 802.5 type LANS), and interoperability is achieved by appropriate frame encapsulation.

A bridge relays individual MAC user data frames between separate MAC protocols of the bridged LAN connected to the ports of the bridge. A MAC entity for each port handles all the media access method dependent functions, i.e., MAC protocol and procedures, as specified in the relevant IEEE 802 standard for that MAC technology. Each bridge port receives and transmits frames to and from the LAN to which it is attached using the services provided by the individual MAC entity associated with that port. Each bridge port also functions as an end station providing MAC service to the LLC layer. All MAC entities communicating across a bridged LAN are uniquely identified by their respective 48-bit MAC addresses. A bridge may use a 48-bit MAC address, or a 16-bit locally administered MAC address. This bridge address must be unique within the extended LAN, and a single unique bridge identifier (ID) is derived from it for the operation of a bridge protocol. Each frame transmitted from a source end station, for example 16, to a destination end station, for example 18, carries the MAC addresses of the end stations respectively in the source and destination address fields of the frame's MAC header. A frame that is to be relayed by every bridge to all its neighboring bridges in an extended LAN contains a bridge group MAC address in the destination address field of the frame's MAC header.

The three basic functions set forth in the present standards of an IEEE 802.1D bridge are:

1) frame forwarding – forward a frame received from one port to another port
2) learning – "learn" and "remember" which port to forward a frame
3) spanning tree algorithm – make sure activated links form no loop, i.e., the bridges and links form a spanning tree

Functions (1) and (2) above are performed with the help of a *Forwarding Database*, or Filtering Database, (see Fig. 7-4 of IEEE 802.ID 1998 Edition), within each bridge. Each bridge keeps a Forwarding Database, hereafter denoted FD, that specifies which port to forward a data frame with a